Efficient Passive Image–Based Hazard Detection for Safe Landing on Mars

Yang Cheng, Larry H. Matthies, Andrew E. Johnson and Aron A. Wolf
Jet Propulsion Laboratory

Introduction

According to NASA's current outline for Mars exploration, a long–range, long duration roving science laboratory is base–lined for the 2007 launch opportunity. This "range–rover" will provide a major improvement in science measurement quantity and quality and will pave the way for a future sample return mission. In order to land such a large payload at promising but difficult–to–reach scientific sites, technology for accurate landing and hazard avoidance must be developed.

Currently, both active sensing and passive sensing approaches for hazard detection are under investigation. Because active sensors (radar and lidar) can directly measure the depth of the sensed terrain, and their data processing algorithms are relatively simple and fast, active sensing for hazard detection is considered the most promising method for the future missions. However, active sensors have their own generic disadvantages; they are often expensive, heavy and power hungry. In contrast, passive sensors (cameras) have lower mass, lower cost, lower power consumption and higher resolution. In addition, camera technology is flight proven and much more mature than active sensing technology; less risk will be introduced into missions when cameras are used. The biggest disadvantage of passive sensing is that the algorithms required for safe landing are often too computationally expensive. In this paper, we describe a novel hazard detection approach using passive imaging. In this approach, algorithmic efficiency has been dramatically improved over existing techniques through the use of onboard inertial measurements, careful timing of sensing and selective processing of image data. Our initial study shows that this new method can reliably detect safe landing sites at the rate needed for Mars landing.

Safe Landing Criteria

Engineering constraints for safe landing will be derived from the 2007 spacecraft design and landing scenario. Since the detailed scenario for the 2007 mission has not yet been planned, we adopted some of the criteria in the guidelines for Mars Surveyor'01 and Mars 05 lander in this study. According to these studies, a safe landing site should satisfy:

- The surface slope must be below 15 degree;
- The probability of landing on a rock greater than 50 cm high should be less than 1%.

Because high–resolution local surface topography is currently unavailable for Mars, hazard detection and avoidance must be done in real–time while the spacecraft is descending. Specifically, the system must be able to detect the hazards in approximately one second.

Approach

In the conventional passive imaging method, hazard detection is accomplished in multiple stages: feature matching, motion adjustment, depth recovery and hazard detection. The conventional algorithms involved in each step are often computationally expensive. However, strategies exist that can be used to speed up this process:

- Motion adjustment can be replaced by measurements from on–board inertial sensors.
- Only a subset of the image that is within the lander control authority needs to be processed.
- The approximation that highly textured image regions correspond to unsafe regions can be employed to mask large parts of the image from expensive slope computations.
- Instead of all feasible landing sites, only the best landing site needs to be determined.

Input into our safe landing site selection algorithm are two images taken as the lander descends to the surface and the change in position and attitude (motion) of the lander between the two images. Our algorithm has two steps: the first step efficiently reduces the amount of data that has to be processed by the more accurate, but computationally expensive, second step.

In the first step the first image is processed to determine multiple candidate landing sites where the image intensity standard deviation in a window around each pixel is minimized. This step is made extremely efficient by utilizing a dynamic programming technique common in stereo matching.

In the second step, the following observations are employed. If the surface terrain corresponding to an image patch is flat, then the change in appearance of the patch during lander motion can be modeled with a homographic transformation. Using the motion between images, the slope of the underlying terrain can be extracted from this homography. Furthermore, if the terrain is not flat, the residual on the estimate of the homography will be large. Using these observations, the procedure during the second step of our safe landing algorithm is as follows. For each candidate landing site, the homography between the window in the first image and the corresponding image patch in the second image is estimated using a nonlinear estimation algorithm. The slope and error residual is extracted from each homography, and the best landing site is selected to as the patch that has the smallest error residual and also satisfies the specified slope constraint.

Results

We have tested the approach on descent images collected in laboratory and near Silver Lake, California. For a 1000x1000 pixel image, the approach takes 0.07 second to select 9 candidate landing sites in the first image, and it takes 1.2 seconds to compute the 9 homography transforms and corresponding slopes. An example the landing site selection process is shown in the figure below. We have also performed a sensitivity analysis that indicates that slope estimation error is a function of three primary factors: the baseline length to spacecraft height ratio, the descent angle and altitude error from other sensors.